

SPECIFICATION AMENDMENTS

Please amend the paragraph beginning on page 9, line 10 of the specification, as follows:

A variety of inband and outband discovery protocols have been developed for obtaining information about data storage network topology and component attributes. With respect to inband component attribute discovery, devices with SCSI interfaces, such as the storage devices 36 and 38, and the gateway 34, can be polled by the inband agent 22 using SCSI queries to obtain device attribute information, including physical and logical storage unit identification information, port information, and the like. The inband agent 22 may also perform self-discovery to obtain attribute information about the managed host 28, such as file system information, operating system information, and the like. Inband topology queries can be performed by the inband agent 22 using the Fibre Channel GS-3 (Generic Services) protocol (FC-GS-3) and the Fibre Channel FS (Framing and Signaling) protocol (FC-FS) to query the switches 30 and 32 and obtain fabric configuration and end node information. In addition to being responsive to queries from the inband agent 22, all devices implementing SCSI interfaces, HBA drivers, and the FC-GS-3 protocol, will generally support the reporting of device and fabric events as they occur in the data storage network.

Please amend the paragraph beginning on page 15, line 12 of the specification, as follows:

The determination of agent capabilities according to step 90, warrants consideration of how often an agent's capabilities are to be determined, and how much capability information is to be obtained. Assuming capability polling is used to determine agent capabilities (as opposed to some other method such as network topology evaluation), the frequency with

which capability polling is conducted needs to be considered. An exemplary schedule would call for a capability poll to be performed after any or all of the following events:

- 1) A storage network event is detected;
- 2) An agent indicates that its capabilities have changed (e.g., it is unable to collect assigned information);
- 3) An agent is added or removed;
- 4) An agent is upgraded.

Fig. 8 illustrates the foregoing capability poll schedule with reference to the network manager, agents and switches of Fig. 4. As shown, a capability poll 90 is conducted in response to a network event 100, a collection failure 102, an agent removal/addition event 104, and an agent upgrade 106.

Please amend the paragraph beginning on page 18, line 19 of the specification, as follows:

The goal of computing agent assignments according to step 92 of Fig. 5 is as follows: Assuming there is a set of agents that can each cover a subset of ($e_1, e_2 \dots e_m$) of discoverable entities, attempt to identify a potentially smaller subset of agents that can cover all the entities in the set of entities ($e_1, e_2 \dots e_m$) along with the assignments for each agent (i.e., the entities that each agent should cover). Where an entity is discoverable by a single agent, that agent will be assigned responsibility for discovering that entity. Where an entity is discoverable by more than one agent, other factors may need to be considered, such as cost (each agent may have a different cost to gather the same information), load (the amount of information that an agent may be assigned to gather), assignment churn (minimizing change in assignments), etc. There are various algorithms that may be used to compute assignments

based on one or more of these factors. One approach, which is set forth by way of example only and not by way of limitation, would seek to implement the following goals:

- 1) Compute a minimum set of agents that can cover all the entities that need to be covered;
- 2) When an entity can be covered by more than one agent with different costs, assign the agent with the lower cost to cover that entity;
- 3) When an entity can be covered by more than one agent with the same cost, assign the agent with lower potential assignments to that entity;
- 4) Compute assignments such that each entity is covered by exactly one agent;
- 5) Attempt to reduce assignment churn irrespective of the order in which the input data arrives;
- 6) Attempt to base assignments on load balancing when all other criteria ~~is~~are the same, recognizing that optimal load balancing may not be possible when performed incrementally as assignments are made. After all initial assignments have been made consider reassignments to optimize load balancing if this can be done without excessive assignment churn.